

Diffusion Monte Carlo: exponential scaling of computational cost for large systems

Norbert Nemec

Department of Physics, University of Cambridge, CB3 0HE, United Kingdom

(Dated: February 11, 2010)

The computational cost of a Monte Carlo algorithm can only be meaningfully discussed when taking into account the magnitude of the resulting statistical error. Aiming for a fixed error per particle, we study the scaling behavior of the diffusion Monte Carlo method for large quantum systems. We identify the correlation within the population of walkers as the dominant scaling factor for large systems. While this factor is negligible for small and medium sized systems that are typically studied, it ultimately shows exponential scaling. The scaling factor can be estimated straightforwardly for each specific system and we find that it is typically only becomes relevant for systems containing more than several hundred atoms.

PACS numbers: 02.70.Ss, 71.15.Nc, 31.15.-p,

I. INTRODUCTION

Today's scientists can choose from a wide range of computational methods for the simulation of quantum mechanical systems. These range from highly efficient semi-empirical methods to density functional methods – offering a practical compromise of efficiency and accuracy – all the way to very accurate quantum chemical methods. Besides these deterministic methods, various stochastic quantum Monte Carlo (QMC) methods are gaining ground, offering exact handling of many strongly correlated systems and scaling up to system sizes that are out of reach for the deterministic competitors.

The two major arguments that are typically brought up in the advocacy of QMC are the excellent parallelizability and the good scaling behavior. Depending on the QMC variant that is chosen, the collection of statistical data points can be performed in parallel with little to no communication, making the method well suited for high performance computers of all architectures. The scaling behavior depends greatly on the details of the system and the method, but it is generally found to be significantly better than that of quantum chemical methods, and linear scaling algorithms have been reported.^{1–8}

A commonly used method for the ab initio simulation of electronic structure is diffusion Monte Carlo (DMC),^{9–11} typically using the fixed node approximation.¹² For this method, the bulk of the computational effort is spent on the repeated evaluation of a trial wave function for electron positions that change step by step, one electron at a time. The trial wave function is usually expressed as a Slater determinant¹³ of single electron orbitals, multiplied by a Jastrow factor¹⁴ to express electron correlations. For single electron orbitals expressed as maximally localized Wannier functions,¹⁵ the local energy can be reevaluated in constant time after a single electron move, leading to an $O(N)$ algorithm for a complete time step of one configuration.^{1–8} As a further refinement to this, trial wave functions for DMC calculations are today commonly expressed in a blip basis,¹⁶ which can be evaluated very efficiently.

In contrast to deterministic methods, however, the

computational cost of a Monte Carlo (MC) simulation is meaningless without specifying the statistical error that is achieved. Deterministic methods typically have systematic errors that are either intrinsic or depend on parameters that do not scale with the system size. The statistical error of MC simulations on the other hand scales very simply with the inverse square root of the run time while the scaling with the system size is a nontrivial issue that depends on details of the method and the system of study. Though the unfavorable scaling of the statistical efficiency of DMC has been demonstrated before,¹⁷ it has – to our knowledge – never been studied systematically.

In this paper, we will present a systematic study of the scaling behavior of QMC calculations aiming for a fixed statistical error bar. The main focus will be on the DMC algorithm including branching and population control as described by Umrigar et al.,¹⁸ other variants




FIG. 1: (color online) Scaling of the effective population size $N_{\text{pop}}^{\text{eff}}$ [see Eq. (8)] in a sample system [$\alpha = 1.5$, see Eq. (13)] with increasing number of atoms N_{atom} . The target population size is fixed, the true population size N_{pop} fluctuates around a slightly lower average (see text). The “error bars” of N_{pop} visualize the increasing population fluctuations $\sigma_{N_{\text{pop}}}$. The effective population drops exponentially, due to increasing correlations within the population.

will be briefly discussed as well. The statements that we will derive are expected to hold for DMC calculations in general, but to simplify understanding, we will consider a “typical” system made up of N similar constituents which we simply call “atoms”. This could be, for example, a crystal in a simulation cell made up of N primitive cells, a cluster of N atoms or a large organic molecule of N comparable groups.

We will begin by deriving several general quantities and continue by demonstrating these in the case of a simple model of N independent hydrogen atoms. From this model we can numerically extract the missing pieces of the scaling behavior, allowing a quantitative estimate of the scaling limit for arbitrary systems. This limit will then be discussed for a number of different sample systems.

II. SCALING OF COMPUTATIONAL COST

The total computational cost of a DMC calculation (optionally split over a number of parallel CPUs) is

$$t_{\text{total}} = N_{\text{step}} \times N_{\text{pop}} \times t_{\text{step}}, \quad (1)$$

where N_{step} is the number of steps in imaginary time, N_{pop} is the average population size and t_{step} is the CPU time needed for one single all-electron step per configuration. Using a so-called “linear scaling” QMC algorithm,¹ each all-electron move scales as

$$t_{\text{step}} \propto N,$$

assuming that the evaluation of the Slater determinant dominates the computational cost. For the moment, we assume that population fluctuations are negligible and N_{pop} can be treated as an external parameter. The influence of the population control will be discussed later on.

The standard error of the total energy can be expressed as

$$\delta E_{\text{total}} = \sqrt{\left(\frac{\tau_{\text{corr}}}{\tau_{\text{step}} N_{\text{step}}}\right) \left(\frac{\chi_{\text{pop}}}{N_{\text{pop}}}\right) \sigma_{\text{dmc}}^2}, \quad (2)$$

with the constituents explained in the following:

The raw DMC variance σ_{dmc}^2 is the variance of the local energy of individual configurations over the whole simulation. Being based on the mixed estimator, σ_{dmc}^2 may deviate from the variance σ_{vmc}^2 of the trial wave function obtained in a VMC run. For typical systems, however, we find that $\sigma_{\text{dmc}}^2 \approx \sigma_{\text{vmc}}^2$. Assuming that the trial wave function of the whole system can be optimized to about the same quality as that of a single constituent N , this variance scales as

$$\sigma_{\text{dmc}}^2 \propto N, \quad (3)$$

since the local energy is dominated by the sum of N independent atomic local energies.

The (integrated) correlation time of a series of data points x_i is given by¹⁹

$$\tau_{\text{corr}} = \tau_{\text{step}} \left(1 + 2 \sum_{j=1}^{\infty} \frac{\langle x_i x_{i+j} \rangle_i - \langle x_i \rangle_i^2}{\langle x_i^2 \rangle_i - \langle x_i \rangle_i^2} \right), \quad (4)$$

where $\langle \cdot \rangle_i$ denotes the arithmetic mean over the index i . τ_{corr} in units of the time step τ_{step} describes the factor by which the number of steps N_{steps} has to be divided to correct the error bar of the result for serial correlation. To obtain $\tau_{\text{corr}}/\tau_{\text{step}}$, as an alternative to computing Eq. (4) directly, one can also use the reblocking method.²⁰ Though the system may have various correlation time scales, some of which depend on the system size, we find that the integrated correlation time that is responsible for the reduction in the resulting accuracy is dominated by the shortest correlation times which depend on local properties, such as the kind of nuclei in the system, but not on the size N .

The population correlation factor $\chi_{\text{pop}} \geq 1$ captures the inefficiency of the process due to population correlation and fluctuation. We will treat this factor as an unknown quantity for the moment and discuss it in detail afterwards.

Using Eqs. (1) and (2) along with the discussed scaling laws, we can express the scaling of the total computational cost as

$$t_{\text{total}} \propto \frac{\chi_{\text{pop}}}{\delta E_{\text{atom}}^2}, \quad (5)$$

so we see that – apart from the factor χ_{pop} – DMC is in fact a constant scaling method if a fixed standard error per atom $\delta E_{\text{atom}} = \delta E_{\text{total}}/N$ is required, as it is the case for example in the study of long-ranged correlations in periodic systems. Of course, memory limitations or implementation issues will limit the size of computable systems. Within these limitations, however, the constant scaling behavior is not so surprising, considering that for collecting statistical data, it does not make any difference whether you simulate N weakly interacting systems in parallel or a single system N times as long. Both result in the same statistical error for the same computational cost.

III. POPULATION CORRELATION

For small enough systems, the factor χ_{pop} is close to one, which may be the reason why, to our knowledge, a systematical study has never been attempted before. When scaling up the system size, however, population correlation becomes important and we need a better understanding of its scaling.

The DMC algorithm is based on a drift-diffusion process with branching and killing of configurations due to fluctuations in the local energy. A freshly branched pair of configurations is identical and thereby fully correlated.

In the following drift-diffusion process, it takes some time to decorrelate, leading to a fluctuating amount of correlation within the population at any time.

We consider a DMC run over N_{step} time steps i . We will first consider a simplified model with constant population of N_{pop} configurations p , each having a local energy E_p^i . A generalization including population fluctuations will follow in the section below.

An effective population size $N_{\text{pop}}^{\text{eff}}$ can be defined as the number of configurations that would result in the same variance of the average as the correlated population. For a long DMC run, the raw DMC variance can be estimated from the averages over all configurations at all time steps as

$$\sigma_{\text{dmc}}^2 = \left\langle \left\langle (E_p^i)^2 \right\rangle_p \right\rangle_i - \left\langle \left\langle E_p^i \right\rangle_p \right\rangle_i^2, \quad (6)$$

while the variance of the population average is defined as

$$\sigma_{\text{pop}}^2 = \left\langle \left\langle E_p^i \right\rangle_p^2 \right\rangle_i - \left\langle \left\langle E_p^i \right\rangle_p \right\rangle_i^2, \quad (7)$$

with the averages abbreviated as $\langle \cdot \rangle_p = \sum_{p=1}^{N_{\text{pop}}} \cdot / N_{\text{pop}}$ and $\langle \cdot \rangle_i = \sum_{i=1}^{N_{\text{step}}} \cdot / N_{\text{step}}$.

In the case of an uncorrelated population, we would find $\sigma_{\text{pop}}^2 = \sigma_{\text{dmc}}^2 / N_{\text{pop}}$, so we can measure the amount of correlation by defining an effective population size as the ratio

$$N_{\text{pop}}^{\text{eff}} = \sigma_{\text{dmc}}^2 / \sigma_{\text{pop}}^2, \quad (8)$$

where $N_{\text{pop}}^{\text{eff}} \leq N_{\text{pop}}$ with equality only in the case of a completely uncorrelated population. By collecting the necessary data during a DMC run, $N_{\text{pop}}^{\text{eff}}$ can be computed at negligible cost. Fig. 1 displays the scaling of the effective population size with increasing system size for a sample system.

Note that the average population N_{pop} is typically slightly lower than the target population $N_{\text{pop}}^{\text{target}}$, because the population control implemented in CASINO uses the linear average $\langle E_p^i \rangle_p$ of the energy instead of the exponential average $\ln \langle \exp(E_p^i) \rangle_p$ which determines the actual growth of the population. Apart from reducing the population size, this has no effect on the statistics or the result.

IV. POPULATION FLUCTUATIONS

Keeping the population size completely fixed as we had assumed in the previous model gives rise to a population control bias. To reduce this bias, the population size N_{pop}^i is allowed to fluctuate and only weakly controlled. The average over all time steps then needs to be weighted. In the simple population control mechanism considered here,¹⁸ the weights are simply defined by the population

size $w_i = N_{\text{pop}}^i$ for each time step i . The resulting total energy average of a DMC run is then

$$E_{\text{tot}} = \frac{1}{\sum_i w_i} \sum_{i=1}^{N_{\text{step}}} w_i \langle E_p^i \rangle_p.$$

In more sophisticated schemes, w_i and N_{pop}^i may be decoupled. To estimate the variance of this weighted average, we can split off the correlation time into a factor and use the estimator of the variance of a weighted average. Viewing the local energies E_p^i as random variables and the weights w_i as constants given by a long DMC run, we can express this as

$$\text{var}[E_{\text{tot}}] = \frac{\tau_{\text{corr}}}{\tau_{\text{step}}} \frac{1}{(\sum_i w_i)^2} \sum_{i=1}^{N_{\text{step}}} w_i^2 \text{var}[\langle E_p^i \rangle_p]. \quad (9)$$

In the same interpretation of E_p^i as random variables, we can replace the averages over time steps in Eqs. (6) and (7) by statistical expectation values $\langle \cdot \rangle$ and write for each single time step

$$\text{var}[\langle E_p^i \rangle_p] = \sigma_{\text{dmc}}^2 - \left\langle \left\langle (E_p^i)^2 \right\rangle_p - \left\langle E_p^i \right\rangle_p^2 \right\rangle.$$

Substituting this into Eq. (9) results in a sum over expectation values, so the $\text{var}[E_{\text{tot}}]$ itself can be written as the expectation value of a single expression which can be expressed as a product

$$\text{var}[E_{\text{tot}}] = \left\langle \frac{\tau_{\text{corr}}}{\tau_{\text{step}}} \times \frac{1}{N_{\text{step}}^{\text{eff}}} \times \frac{1}{N_{\text{pop}}^{\text{eff}}} \times \sigma_{\text{dmc}}^2 \right\rangle,$$

with generalized expressions for the effective step number and population size

$$\frac{1}{N_{\text{step}}^{\text{eff}}} = \frac{\sum_i (w_i^2)}{(\sum_i w_i)^2}$$

$$\frac{1}{N_{\text{pop}}^{\text{eff}}} = 1 - \frac{1}{\sigma_{\text{dmc}}^2} \frac{\sum_i w_i^2 \left(\left\langle (E_p^i)^2 \right\rangle_{\text{pop}} - \left\langle E_p^i \right\rangle_{\text{pop}}^2 \right)}{\sum_i w_i^2},$$

where the quadratic appearance of the weights makes the effective population size sensitive to population fluctuations as well.

To estimate the statistical efficiency of the DMC algorithm, it is most useful to combine both quantities into the definition

$$\chi_{\text{pop}} = \frac{N_{\text{step}} \times N_{\text{pop}}}{N_{\text{step}}^{\text{eff}} \times N_{\text{pop}}^{\text{eff}}}. \quad (10)$$

For a DMC run that is sufficiently long that the set of weights w_i is a good representation of the statistical distribution, the quantity χ_{pop} is an unbiased estimator. It corresponds exactly to the quantity in Eq. (2) and remains directly proportional to the total CPU cost according to Eq. (5).

V. ASYMPTOTICS OF χ_{pop}

The asymptotic behavior of χ_{pop} for weak population correlation can be derived by a few simple arguments. Assume, for a moment, the local energy E_{loc} of configurations to be normally distributed as

$$p(E_{\text{loc}}) = \frac{1}{\sigma_{\text{dmc}}\sqrt{2\pi}} \exp\left(-\frac{E_{\text{loc}}^2}{2\sigma_{\text{dmc}}^2}\right).$$

A single configuration with $E_{\text{loc}} < 0$ will branch at a rate of $-E_{\text{loc}}$. Integrated over the distribution of E_{loc} , this leads to a branching rate per configuration of

$$\begin{aligned} \tau_{\text{branch}}^{-1} &= \int_{-\infty}^0 dE_{\text{loc}} p(E_{\text{loc}}) (-E_{\text{loc}}) \\ &= \frac{\sigma_{\text{dmc}}}{\sqrt{2\pi}}. \end{aligned}$$

It is, of course, known that the true distribution of the local energy is far from normal.²¹ However, we can significantly relax the previous assumption, considering that we essentially obtained the ratio between standard deviation and mean absolute deviation which holds approximately for a wide range of distributions.²²

Starting from an initially uncorrelated population of size N_{pop} , the population after branching is $N_{\text{pop}} + 1$ with two identical configurations. The population mean is equivalent to that over $N_{\text{pop}} - 1$ correlations of single weight and one of double weight which has the variance

$$(\sigma_{\text{pop}}^2)' = \sigma_{\text{dmc}}^2 \frac{(N_{\text{pop}} - 1) \times 1^2 + 1 \times 2^2}{((N_{\text{pop}} - 1) \times 1 + 1 \times 2)^2}.$$

For $N_{\text{pop}} \gg 1$, the effective population size after the branching is therefore $(N_{\text{pop}}^{\text{eff}})' = N_{\text{pop}} - 1$.

To keep the population stable, the branching and killing rates have to be equal. For a weakly correlated population, it is trivial to see that a killing event reduces the (effective) population by 1 as well.

After branching, the two copies evolving independently take an effective time of $\tau_{\text{corr}}/2$ to decorrelate. (Since τ_{corr} measures the amount of statistical data that is lost due to serial correlation, which is exactly the quantity that we want to measure for population correlation as well.)

For $\tau_{\text{branch}} \gg \tau_{\text{corr}}$, the effective population is reduced by 1 at a rate of $2N_{\text{pop}}\tau_{\text{branch}}^{-1}$ and restored to N_{pop} within $\tau_{\text{corr}}/2$. On average this gives

$$N_{\text{pop}}^{\text{eff}} = N_{\text{pop}} - 2N_{\text{pop}}\tau_{\text{branch}}^{-1}\tau_{\text{corr}}/2$$

or, since population fluctuations can be neglected,

$$(\chi_{\text{pop}})_{\tau_{\text{corr}}\sigma_{\text{dmc}} \rightarrow 0} \rightarrow 1 + \frac{\tau_{\text{corr}}\sigma_{\text{dmc}}}{\sqrt{2\pi}}. \quad (11)$$

VI. MODEL DMC PROCESS

To study the dependence of χ_{pop} on the system parameters beyond the perturbative regime, we have implemented the full DMC algorithm on top of a minimal model of a correlated diffusion process. Each configuration is reduced to a single random variable with a simple exponential autocorrelation so that σ_{dmc} and τ_{corr} are free parameters. The value of the variable is used directly as the local energy for the branching process.

The population control mechanism described by Umrigar et al. introduces an additional parameter τ_{ceref} with the dimension of time (implemented in CASINO as parameter `cerefdmc`, used in updating the reference energy E_{ref} ; $\tau_{\text{ceref}}/\tau_{\text{step}}$ corresponds to g in the original publication¹⁸). To avoid frequent population instabilities for extreme parameter settings, we have restricted the population size to a window around the target population and decoupled the total weight from the total population size outside of this window, allowing recovery from explosions or starvation without introducing any additional bias. All the results presented below are in the regime where this mechanism has no significant impact on the efficiency.

With each run, one obtains the factor χ_{pop} as a function of the parameters σ_{dmc} , τ_{corr} , τ_{ceref} and τ_{step} . The result must be dimensionless, reducing the number of relevant parameters by one. Furthermore, one is interested in the limit $\tau_{\text{step}} \rightarrow 0$. We find that for $\tau_{\text{step}} \lesssim 0.1 \times \min(\tau_{\text{corr}}, \tau_{\text{ceref}}, 1/\sigma_{\text{dmc}})$, the inefficiency factor χ_{pop} becomes practically independent of τ_{step} in all cases. We can combine the remaining free parameters into $\sigma_{\text{dmc}}\tau_{\text{corr}}$ and $\tau_{\text{ceref}}/\tau_{\text{corr}}$, leading to the results displayed in Fig. 2.

Most significantly, we find that Eq. (11) is not only confirmed in the perturbative limit, but its exponential continuation gives a strict lower limit for the inefficiency factor

$$\chi_{\text{pop}} \geq \exp\left(\sigma_{\text{dmc}}\tau_{\text{corr}}/\sqrt{2\pi}\right) \quad (12)$$

where the deviation from this exponential depends on the ratio $\tau_{\text{ceref}}/\tau_{\text{corr}}$.

VII. HYDROGEN SAMPLE SYSTEM

To demonstrate our result in a real calculation, we have performed various DMC runs using the CASINO program.²⁵ We chose a system of N hydrogen atoms placed several thousand atomic units apart to make them effectively independent. As a trial wave function, we used the exact ground state with a detuning parameter α and an additional term to satisfy the Kato cusp condition,²⁶ centered on each hydrogen atom

$$\Psi_{\alpha}(r) = \alpha e^{-\alpha r} + (1 - \alpha) e^{-(\alpha+1)r}. \quad (13)$$




FIG. 2: (color online) Inefficiency factor χ_{pop} computed for a model DMC process (see text) in dependence of the two relevant parameters $\sigma_{\text{dmc}}\tau_{\text{corr}}$ and $\tau_{\text{ceref}}/\tau_{\text{corr}}$. The exponential law extrapolated from the perturbative limit is found to give a strict lower limit for χ_{pop} . Solid circles refer to standard DMC with branching. Hollow circles refer to minimal stochastic reconfiguration MC (SRMC).^{23,24}

We performed a large variety of runs on this model system with system sizes $N \in \{1, 2, 4, 8, \dots, 64\}$, and detuning parameters $\alpha \in \{1.1, \dots, 3.0\}$ and target population $N_{\text{pop}} = 200$. The DMC time step was set to $\tau_{\text{step}} = 0.02$ in all cases.

The variance σ_{dmc}^2 was found to be equal to σ_{vmc}^2 within the statistical error in all cases. In each case, the population correlation factor χ_{pop} was determined from the variances using Eq. (10).

Obtaining a precise value for the correlation time τ_{corr} takes an extremely large amount of data in either of the two methods described above. For a reasonably precise value, we performed a very long DMC run ($N_{\text{step}} > 10^7$) on a single atom for each value of α . Several tests on larger systems confirmed that the same value holds for larger numbers of atoms N . Since τ_{corr} is independent of τ_{step} only if $\tau_{\text{corr}} \gg \tau_{\text{step}}$, we performed these runs for the same step size as the main calculations.

As in Fig. 2, we plot the population correlation factor χ_{pop} from all our calculations as a function of the product $\sigma_{\text{dmc}}\tau_{\text{corr}}$ and again find the exponential lower bound described by Eq. (12), as displayed in Fig. 3.

VIII. ANALYSIS OF VARIOUS SAMPLE SYSTEMS

The exponential law in Eq. (12) has severe implications for the scaling of the DMC method. Following Eqs. (3)




FIG. 3: (color online) Data from many different calculations on a model system of N independent hydrogen atoms with a detunable trial wave function [Eq. (13)]. σ_{dmc} and χ_{pop} were directly obtained from each run. τ_{corr} was determined from a single, very long run for each type of atomic wave function. Population control is kept at the default $t_{\text{ceref}} = 1$.

and (5), the total CPU cost becomes

$$t_{\text{total}} \propto \frac{\exp(X\sqrt{N})}{\delta E_{\text{atom}}^2}$$

or worse. So, even if population correlation may not be an issue yet for most applications, it will eventually lead to an exponential scaling of the cost. The factor X can be reduced by optimizing the wave function, but the gain that is possible with reasonable effort is very limited.

Table I lists a selection of sample systems showing the size at which the exponential scaling becomes observable. The integrated correlation time τ_{corr} [via Eq. (4)] and the raw variance σ_{dmc}^2 were computed for very small systems and Eq. (12) was then used to estimate the size at which a comparable system would show significant population correlation. All values should be understood as rough estimates. The trial wave functions were either taken from a library of examples or optimized with moderate effort. Further optimizations could certainly reduce σ_{dmc}^2 and thereby shift the onset of significant population correlation. Typically, however, significant effort is necessary even for minor improvements using optimizations beyond the standard Jastrow terms.

IX. ALTERNATIVE VARIANTS OF QMC

To this point the discussion was centered on the conventional DMC algorithm including drift and branching. In the following, we will briefly discuss a number of alternative QMC algorithms in view of the population correlation scaling.

First, it is clear that population correlation can only be caused by some form of branching. The variational MC

atoms (ae)	τ_{corr}	$\sigma_{\text{dmc}}^2/\text{atm}$	$\chi_{\text{pop}} = 2$
He	0.5	0.0044	2700 atoms
C	0.4	0.16	140 atoms
Ar	0.04	8.0	250 atoms
molecules (ae)		$\sigma_{\text{dmc}}^2/\text{mlc}$	
H ₂ O	0.1	0.58	550 molec.
CH ₄	0.3	0.24	120 molec.
C ₂ H ₄	0.4	0.51	38 molec.
SO ₂	0.06	7.5	105 molec.
crystals		$\sigma_{\text{dmc}}^2/\text{atom}$	
diamond (pp)	0.15	0.23	630 atoms
diamond (ae)	0.1	2.3	133 atoms
graphite (pp)	0.3	0.20	135 atoms
silicon (pp)	0.4	0.052	328 atoms
electron gas		$\sigma_{\text{dmc}}^2/\text{elec.}$	
3d crystal ($r_s = 1$)	0.2	0.26	193 elec.
3d fluid ($r_s = 5$)	5	4.2×10^{-4}	330 elec.
3d fluid ($r_s = 10$)	16	5.1×10^{-5}	242 elec.
2d crystal ($r_s = 1$)	0.4	0.038	570 elec.
2d fluid ($r_s = 1$)	0.3	0.033	1154 elec.

TABLE I: Estimated values for various sample systems. The last column gives the system size based on Eq. (12) at which the population correlation becomes significant with $\chi_{\text{pop}} = 2$. Beyond this size, the DMC method must be expected to become exponentially inefficient. The first three categories are based on either all-electron (ae) or pseudopotential (pp) wave functions with optimized Jastrow factors. All numbers should be understood as rough estimates based on moderately optimized trial wave functions. Reducing σ_{dmc}^2 by further optimization will shift the onset of the inefficiency by the same factor.

(VMC) algorithm, which samples an explicitly known wave function, clearly does not have this feature. A variant of DMC with branching switched off (sometimes referred to as “pure” DMC) also features a completely uncorrelated population. If all configurations are fixed to the same statistical weight, this process produces the same distribution of configurations and thereby the same total energy as VMC and can therefore be seen as a variant of the former.

If, on the other hand, each configuration in a pure DMC run carries a statistical weight evolving with the fluctuations in the local energy, the effective population size $N_{\text{pop}}^{\text{eff}}$ is reduced in the same way as it would be when branching were allowed, with the only difference that decorrelation does not happen and the method becomes exponentially unstable with simulation time.²⁴

A number of variants of the DMC algorithms keep the population size fixed and include branching in form of stochastic reconfiguration, duplicating some configurations and deleting others.^{23,24,27} While the population is fixed, the total weight is allowed to fluctuate independently and the population control is replaced by a weight control mechanism.

Our definition of the population correlation factor χ_{pop} in Eq. (10) is already kept general enough to capture the effects of weight fluctuations within the population along the correlations within the population and to capture the

fluctuations of the total weight along with the population fluctuations. Tests on several variants of the branching strategy confirmed that these have no influence on the exponential lower bound of χ_{pop} but only affect how far the actually measured χ_{pop} exceeds the predicted exponential scaling (see Fig. 2).

One remaining option to limit χ_{pop} is the use of strong population control on a small population, accepting a significant population bias. The extreme case of this strategy would be a single walker with the weight renormalized after every step, leading exactly to the VMC distribution. We can, therefore, tune between exponentially scaling statistical inefficiency and a population bias that ultimately leads to recovering the VMC algorithm, which - as we know - does not suffer from exponential scaling.

Variants of QMC such as path integral MC (PIMC)²⁸ or reptation Monte Carlo (RMC)²⁹ are somewhat related to DMC in the sense that they are based on a drift-diffusion process in imaginary time. Unlike DMC, however, these methods are based on a true Metropolis algorithm without the need for branching. Population correlation does not occur and the statistical weight fluctuations are not a problem. Instead, an analysis of the statistical efficiency of these methods would need to take into account the correlation time and its scaling with system size.

X. CONCLUSIONS

To conclude, we have derived an expression for the scaling behavior of DMC calculations when aiming at a fixed statistical precision per particle. Using a linear scaling algorithm for an individual time step, constant scaling of the total computational cost for the energy per particle is possible in principle, except for a factor χ_{pop} , which quantifies the correlation within the population of walkers. The exact value of χ_{pop} was derived in the perturbative limit, depending only on the correlation time and the raw variance of the DMC process. Based on numerical evidence, we demonstrated that an exponential extrapolation of the perturbative law gives a strict lower bound to the inefficiency factor χ_{pop} . From this, it follows that the DMC algorithm generally scales at least exponentially in the square root of the system size.

The numbers for actual sample systems indicate that this exponential scaling should not even be observable in most DMC based studies done so far, leaving plenty of room to do interesting research with the DMC method.

Alternative schemes for branching and population control that have been suggested^{23,24,27} may certainly influence the efficiency of the algorithm. The exponential lower bound of the statistical inefficiency, however, may at best be shifted over towards an exponentially scaling population control bias.³⁰

It must be stressed that this exponential scaling factor is specific to the DMC method and does not occur in other methods like VMC. It is not linked to the more

fundamental fermion sign problem³¹ and it is not limited to certain observables.³²

In fact, the exponential scaling may be a symptom of the very nature of the DMC process. In general, Markov-chain MC methods such as VMC exhibit excellent scaling behavior. DMC however, is not based on a Markov process but rather on the simulation of a time dependent stochastic diffusion process. As such, it must be expected to suffer from the exponential accumulation of errors inherent in the simulation of time evolution in non-integrable systems. The population control that is necessary to stabilize the process might then necessarily lead to exponential scaling either in the bias or the efficiency of the process. The only way to overcome this problem

might then be to resort to alternative QMC methods like VMC, PIMC or RMC that are based the stochastic computation of a multi-dimensional integral in the original spirit of Markov-chain MC methods.

XI. ACKNOWLEDGMENTS

We acknowledge fruitful discussions with Richard Needs, Neil Drummond and Matthew Foulkes. This work was funded by the DAAD and the EPSRC. The computations were performed using the facilities of the University of Cambridge High Performance Computing Service.

-
- ¹ A. J. Williamson, R. Q. Hood, and J. C. Grossman, Phys. Rev. Lett. **87**, 246406 (2001).
 - ² S. Manten and A. Lüchow, in *Recent Advances in Quantum Monte Carlo Methods*, edited by J. William A. Lester, S. M. Rothstein, and S. Tanaka (World Publishing, 2002), vol. II.
 - ³ S. Manten and A. Lüchow, J. Chem. Phys. **119**, 1307 (2003).
 - ⁴ D. Alfè and M. J. Gillan, J. Phys.: Condens. Matter **16**, L305 (2004).
 - ⁵ F. A. Reboredo and A. J. Williamson, Phys. Rev. B **71**, 121105(R) (2005).
 - ⁶ A. Aspuru-Guzik, R. Salomón-Ferrer, B. Austin, and J. William A. Lester, J. Comput. Chem. **26**, 708 (2005).
 - ⁷ J. Kussmann and C. Ochsenfeld, J. Chem. Phys. **128**, 134104 (2008).
 - ⁸ P. K. V. V. Nukala and P. R. C. Kent, J. Chem. Phys. **130**, 204105 (2009).
 - ⁹ J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975).
 - ¹⁰ D. M. Ceperley and B. J. Alder, Phys. Rev. Lett. **45**, 566 (1980).
 - ¹¹ W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).
 - ¹² P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, Jr., J. Chem. Phys. **77**, 5593 (1982).
 - ¹³ J. C. Slater, Phys. Rev. **34**, 1293 (1929).
 - ¹⁴ R. Jastrow, Phys. Rev. **98**, 1479 (1955).
 - ¹⁵ G. H. Wannier, Phys. Rev. **52**, 191 (1937).
 - ¹⁶ D. Alfè and M. J. Gillan, Phys. Rev. B **70**, 161101(R) (2004), cond-mat/0407037.
 - ¹⁷ S. Moroni (2008), presentation Feb 2008.
 - ¹⁸ C. J. Umrigar, M. P. Nightingale, and K. J. Runge, J. Chem. Phys. **99**, 2865 (1993).
 - ¹⁹ U. Wolff, Comput. Phys. Commun. **156**, 143 (2004).
 - ²⁰ H. Flyvbjerg and H. G. Petersen, J. Chem. Phys. **91**, 461 (1989).
 - ²¹ J. R. Trail, Phys. Rev. E **77**, 016703 (2008).
 - ²² S. Gorard, British J. Educational Studies **53**, 417 (2005).
 - ²³ M. Calandra Buonaurea and S. Sorella, Phys. Rev. B **57**, 11446 (1998).
 - ²⁴ R. Assaraf, M. Caffarel, and A. Khelif, Phys. Rev. E **61**, 4566 (2000).
 - ²⁵ R. J. Needs, M. D. Towler, N. D. Drummond, and P. López Ríos, *CASINO user's guide, version 2.3.0* (University of Cambridge, U.K., 2008), URL <http://www.tcm.phy.cam.ac.uk/~mdt26/casino2.html>.
 - ²⁶ T. Kato, Commun. Pure Appl. Math. **10**, 151 (1957).
 - ²⁷ A. Jones, A. Thompson, J. Crain, M. H. Müser, and G. J. Martyna, Phys. Rev. B **79**, 144119 (2009).
 - ²⁸ D. M. Ceperley, Rev. Mod. Phys. **67**, 279 (1995).
 - ²⁹ S. Baroni and S. Moroni, Phys. Rev. Lett. **82**, 4745 (1999).
 - ³⁰ N. D. Drummond, R. J. Needs, A. Sorouri, and W. M. C. Foulkes, Phys. Rev. B **78**, 125106 (2008), 0806.0957.
 - ³¹ M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).
 - ³² G. L. Warren and R. J. Hinde, Phys. Rev. E **73**, 056706 (2006).